Analyzing sets of packet classification rules in order to optimize packet classification based on DCFL

Michal Kekely
xkekel01@stud.fit.vutbr.cz

October 1, 2016

This work focuses on the topic of analyzing packet classification rules. Typical use case of packet classification is in various firewalls, network monitoring devices, probes used for lawful interception of network communication and even in network switches and routers. One of the aproaches used to keep up with increasing speed requirements imposed on those devices is to decompose the packet classification into more independent parts. Individual parts can then be computed in parallel, which is perfectly aligned with the use of FPGA and ASIC based architectures. Examples of such decompositional algorithms are DCFL [1], HiCuts [2] and HyperCuts [3]. However, performance of all of those algorithms is dependent on the actual set of classification rules, with worst-case scenarios being orders of magnitude slower than best-case scenarios. The decrease in performance is caused mainly by the fact that some rules are nested within each other (meaning that for example IP address in one of the rules is prefix of an IP adress in another or that range of ports covered by one rule is obtained in range of ports covered by another rule).

Current optimizations of DCFL algorithm, such as the ones mentioned in [4], are focused mainly on optimizing the architecture. However, there have been no attempts to analyze and understand the sets of rules and relations between individual rules that cause the decrease in performance of DCFL algorithm. Focusing on classification rules has potential to yield great results. For example if we are able to identify subsets of rules that are causing the decrease in performance, we can just exclude some of those subsets from being processed by the main DCFL algorithm and process them in another way (e.g. by using small ternary content adressable memory running in parallel). Alternatively we can use two DCFL engines running in parallel and just split rules from the subsets between them.

Main focus of this work is to find a suitable mathematical model that describes classification rules and relations contained within them. It is important that it would also be possible to feasibly map classification rules onto such a model. Additionaly it should be possible to identify and optimize potentionaly problematic sets of rules using said model.

References:

[1] Taylor, D. E.; Turner, J. S.: Scalable Packet Classification using Distributed Crossproducting of Field Labels. In *IEEE INFOCOM 2005*, 2005.

[2] Gupta, P.; McKeown, N.: Packet Classification using Hierarchical Intelligent Cuttings. 1999.

[3] Singh, S.; Baboescu, F.; Varghese, G.: Packet Classification Using Multidimensional Cutting. In *Proceedings of the 2003 Conference on Applications,Technologies, Architectures, and Protocols for Computer Communications*, SIGCOMM '03, New York, NY, USA: ACM, 2003.

[4] Sun, L.; Le, H.; Prasanna, V., K.: Optimizing Decomposition-based Packet Classification Implementation on FPGAs, 2011.